

Dynamics of free group automorphisms

Peter Brinkmann

June 17, 2008

Abstract

We present a coarse convexity result for the dynamics of free group automorphisms: Given an automorphism ϕ of a finitely generated free group F , we show that for all $x \in F$ and $0 \leq i \leq N$, the length of $\phi^i(x)$ is bounded above by a constant multiple of the sum of the lengths of x and $\phi^N(x)$, with the constant depending only on ϕ .

Introduction

The following theorem is the main result of this paper. It follows from a technical result (Theorem 1.9) that uses the machinery of improved relative train track maps of Bestvina, Feighn, and Handel [BFH00].

Theorem 0.1. *Let $\phi: F \rightarrow F$ be an automorphism of a finitely generated free group. Then there exists a constant $K \geq 1$ such that for any pair of exponents N, i satisfying $0 \leq i \leq N$, the following two statements hold:*

1. *If w is a cyclic word in G , then*

$$\|\phi_{\#}^i(w)\| \leq K (\|w\| + \|\phi_{\#}^N(w)\|),$$

where $\|w\|$ is the length of the cyclic reduction of w with respect to some word metric on F .

2. *If w is a word in F , then*

$$|\phi_{\#}^i(w)| \leq K (|w| + |\phi^N(w)|),$$

where $|w|$ is the length of w .

Given an improved relative train track representative of some power of ϕ , the constant K can be computed.

Remark 0.2 (A note on computability). Given an automorphism $\phi: F \rightarrow F$, we can compute a relative train track representative of ϕ [BH92, DV96]. The construction of *improved* relative train track maps, however, involves a compactness argument in a universal cover [BFH00, Proof of Proposition 5.4.3] that is not

constructive. A number of algorithmic improvements of relative train tracks appear in [Bri07], in the context of an algorithm that detects automorphic orbits in free groups.

The statement of the theorem does not depend on the choice of generators of F . The intuitive meaning of the theorem is that the map $i \mapsto |\phi^i(w)|$ is coarsely convex for all words $w \in F$. Klaus Johannson informed me that a similar result is a folk theorem in the case of surface homeomorphisms. Also, while free-by-cyclic groups are not, in general, CAT(0)-groups [Ger94], Theorem 0.1 suggests that their dynamics mimics that of CAT(0)-groups. Theorem 0.1 complements the following strong convexity result in an important special case.

Theorem 0.3 ([Bri00]). *If $\phi: F \rightarrow F$ is an atoroidal automorphism, i.e., ϕ has no nontrivial periodic conjugacy classes, then ϕ is hyperbolic, i.e., there exists a constant $\lambda > 1$ such that*

$$|x| \leq \lambda \max \{ |\phi^{\pm 1}(x)| \}$$

for all $x \in F$.

I originally set out to prove Theorem 0.1 because it immediately implies that in a free-by-cyclic group

$$\Gamma = F \rtimes_{\phi} \mathbb{Z} = \langle x_1, \dots, x_n, t \mid t^{-1}x_it = \phi(x_i) \rangle,$$

words of the form $t^{-k}wt^k\phi^k(w^{-1})$ satisfy a quadratic isoperimetric inequality. (Note, however, that Theorem 0.1 is stronger than the mere existence of a quadratic isoperimetric inequality for such words.) Natasa Macura previously proved a quadratic isoperimetric inequality for mapping tori of automorphisms of polynomial growth [Mac00]. Martin Bridson and Daniel Groves have since proved that all free-by-cyclic groups satisfy a quadratic isoperimetric inequality [BG]. They also obtain a new proof of Theorem 0.1 as an application of their techniques.

In Section 1, we review the pertinent definitions and results regarding train track maps from [BFH00]. We also state the main technical result, Theorem 1.9, and we show how Theorem 0.1 follows from Theorem 1.9. Section 2 provides some more results on train tracks and automorphisms of free groups. Section 3 introduces some notation and terminology and lists a number of examples that illustrate some of the issues and subtleties that need to be addressed in the proof of Theorem 1.9. Section 4 establishes a technical proposition that may be of independent interest. Finally, Section 5 and Section 6 contain the proof of Theorem 1.9.

I would like to express my gratitude to Ilya Kapovich for many helpful discussions, to Mladen Bestvina for patiently answering my questions, to Steve Gersten for encouraging me to write up this result for its own sake, to the University of Osnabrück for their hospitality, and to Swarup Gadde and the University of Melbourne as well as the Max-Planck-Institute of Mathematics for their hospitality and financial support. Klaus Johannson and Richard Weidmann kindly served as a sounding board while I was working on the exposition of this paper.

1 Improved relative train track maps

In this section, we review the theory of train tracks developed in [BH92, BFH00]. We will restrict our attention to the collection of those results that we will use in this paper.

Given an automorphism $\phi \in \text{Aut}(F)$, we can find a based homotopy equivalence $f: G \rightarrow G$ of a finite connected graph G such that $\pi_1(G) \cong F$ and f induces ϕ . This observation allows us to apply topological techniques to automorphisms of free groups. In many cases, it is convenient to work with outer automorphisms. Topologically, this means that we work with homotopy equivalences rather than based homotopy equivalences.

Oftentimes, a homotopy equivalence $f: G \rightarrow G$ will respect a *filtration* of G , i. e., there exist subgraphs $G_0 = \emptyset \subset G_1 \subset \dots \subset G_k = G$ such that for each filtration element G_r , the restriction of f to G_r is a homotopy equivalence of G_r . The subgraph $H_r = \overline{G_r \setminus G_{r-1}}$ is called the *r-th stratum* of the filtration. We say that a path ρ has *nontrivial intersection* with a stratum H_r if ρ crosses at least one edge in H_r .

If E_1, \dots, E_m is the collection of edges in some stratum H_r , the *transition matrix* of H_r is the nonnegative $m \times m$ -matrix M_r whose ij -th entry is the number of times the f -image of E_j crosses E_i , regardless of orientation. M_r is said to be *irreducible* if for every tuple $1 \leq i, j \leq m$, there exists some exponent $n > 0$ such that the ij -th entry of M_r^n is nonzero. If M_r is irreducible, then it has a maximal real eigenvalue $\lambda_r \geq 1$ [Gan59]. We call λ_r the *growth rate* of H_r .

Given a homotopy equivalence $f: G \rightarrow G$, we can always find a filtration of G such that each transition matrix is either a zero matrix or irreducible. A stratum H_r in such a filtration is called *zero stratum* if $M_r = 0$. H_r is called *exponentially growing* if M_r is irreducible with $\lambda_r > 1$, and it is called *polynomially growing* if M_r is irreducible with $\lambda_r = 1$.

An unordered pair of edges in G originating from the same vertex is called a *turn*. A turn is called *degenerate* if the two edges are equal. We define a map $Df: \{\text{turns in } G\} \rightarrow \{\text{turns in } G\}$ by sending each edge in a turn to the first edge in its image under f . A turn is called *illegal* if its image under some iterate of Df is degenerate, *legal* otherwise.

An edge path $\rho = E_1 E_2 \dots E_s$ is said to contain the turns (E_i^{-1}, E_{i+1}) for $1 \leq i < s$. ρ is said to be *legal* if all its turns are legal, and a path $\rho \subset G_r$ is *r-legal* if no illegal turn in ρ involves an edge in H_r .

Let ρ be a path in G . In general, the composition $f^k \circ \rho$ is not an immersion, but there is exactly one immersion that is homotopic to $f^k \circ \rho$ relative endpoints. We denote this immersion by $f_{\#}^k(\rho)$, and we say that we obtain $f_{\#}^k(\rho)$ from $f^k \circ \rho$ by *tightening*. If σ is a circuit in G , then $f_{\#}^k(\sigma)$ is the immersed circuit homotopic to $f^k \circ \sigma$.

Remark 1.1. A path is tightened by cancelling adjacent pairs of inverse edges until no inverse pairs are left. The result of such a sequence of cancellations is uniquely determined, but the sequence is not. For instance, $EE^{-1}E$ may be

tightened as $E(E^{-1}E)$ or $(EE^{-1})E$.

Convention 1.2. Let $\rho_i, i = 1, \dots, k$ be paths that can be concatenated to form a path $\rho = \rho_1 \rho_2 \cdots \rho_k$. When tightening $f(\rho)$ to obtain $f_{\#}(\rho)$, we adopt the convention that we first tighten the images of ρ_i to $f_{\#}(\rho_i)$. In a second step, we tighten the concatenation $f_{\#}(\rho_1) \cdots f_{\#}(\rho_k)$ to $f_{\#}(\rho)$.

In many situations, the length of a subpath ρ_i will be greater than the number of edges that cancel at either end, in which case it makes sense to talk about edges in $f_{\#}(\rho)$ *originating from* ρ_i .

A path ρ is a *(periodic) Nielsen path* if $f_{\#}^k(\rho) = \rho$ for some $k > 0$. In this case, the smallest such k is the *period* of ρ . A Nielsen path ρ is called *indivisible* if it cannot be expressed as the concatenation of shorter Nielsen paths. A path ρ is a *pre-Nielsen path* if $f_{\#}^k(\rho)$ is Nielsen for some $k \geq 0$.

A decomposition of a path $\rho = \rho_1 \cdot \rho_2 \cdots \rho_s$ into subpaths is called a *k-splitting* if $f_{\#}^k(\rho) = f_{\#}^k(\rho_1) \cdots f_{\#}^k(\rho_s)$. Such a decomposition is a *splitting* if it is a *k-splitting* for all $k > 0$. We will also use the notion of *k-splittings* of circuits $\sigma = \rho_1 \cdot \rho_2 \cdots \rho_s$, which requires, in addition, that there be no cancellation between $f_{\#}^k(\rho_s)$ and $f_{\#}^k(\rho_1)$.

The following theorem was proved in [BH92].

Theorem 1.3 ([BH92, Theorem 5.12]). *Every outer automorphism \mathcal{O} of F is represented by a homotopy equivalence $f: G \rightarrow G$ such that each exponentially growing stratum H_r has the following properties:*

1. *If E is an edge in H_r , then the first and last edges in $f(E)$ are contained in H_r .*
2. *If β is a nontrivial path in G_{r-1} with endpoints in $G_{r-1} \cap H_r$, then $f_{\#}(\beta)$ is nontrivial.*
3. *If ρ is an r -legal path, then $f_{\#}(\rho)$ is an r -legal path.*

We call f a *relative train track map*.

A path ρ in G is said to be of *height r* if $\rho \subset G_r$ and $\rho \not\subset G_{r-1}$. If $H_r = \{E_r\}$ is a polynomially growing stratum, then *basic paths* of height r are of the form $E_r \gamma$ or $E_r \gamma E_r^{-1}$, where γ is a path in G_{r-1} . If τ is a closed Nielsen path in G_{r-1} and $f(E_r) = E_r \tau^l$ for some $l \in \mathbb{Z}$, then paths of the form $E_r \tau^k$ and $E_r \tau^k E_r^{-1}$ are *exceptional paths* of height r . Moreover, if $s < r$, $\tau \subset G_{s-1}$, and $f(E_s) = E_s \tau^m$, then $E_r \tau^k E_s^{-1}$ is also an exceptional path of height r .

For our purposes, the properties of relative train track maps are not strong enough, so we will use the notion of improved train track maps constructed in [BFH00]. We only list the properties used in this paper.

Theorem 1.4 ([BFH00, Theorem 5.1.5, Lemma 5.1.7, and Proposition 5.4.3]). *For every outer automorphism \mathcal{O} of F , there exists an exponent $k > 0$ such that \mathcal{O}^k is represented by a relative train track map $f: G \rightarrow G$ with the following additional properties:*

1. If H_r is a zero stratum, then H_{r+1} is an exponentially growing stratum, and the restriction of f to H_r is an immersion. H_r is a zero stratum if and only if it is the union of the contractible components of G_r .
2. If v is a vertex, then $f(v)$ is a fixed vertex. If H_r is a polynomially growing stratum and G' is the collection of noncontractible components of G_{r-1} , then all vertices in $H_r \cap G'$ are fixed.
3. If H_r is an exponentially growing stratum, then there is at most one indivisible Nielsen path τ of height r . If τ is not closed and if it starts and ends at vertices, then at least one endpoint of τ is not contained in $H_r \cap G_{r-1}$.
4. If H_r is a polynomially growing stratum, then H_r consists of a single edge E_r , and $f(E_r) = E_r \cdot u_r$ for some closed path $u_r \subset G_{r-1}$ whose base point is fixed by f .

If $\sigma \subset G_r$ is a basic path of height r that does not split as a concatenation of two basic paths of height r or as a concatenation of a basic path of height r with a path contained in G_{r-1} , then either $f_{\#}^k(\sigma) = E_r \cdot \sigma'$ for some $k \geq 0$, or u_r is a Nielsen path and $f_{\#}^k(\sigma)$ is an exceptional path of height r for some $k \geq 0$.

We call f an *improved relative train track map*.

Finally, we state a lemma from [BFH00] that simplifies the study of paths intersecting strata of polynomial growth.

Lemma 1.5 ([BFH00, Lemma 4.1.4]). *Let $f: G \rightarrow G$ be an improved train track map with a polynomially growing stratum H_r . If ρ is a path in G_r , then it splits as a concatenation of basic paths of height r and paths in G_{r-1} .*

Remark 1.6. In fact, part 4 of Theorem 1.4 implies that subdividing ρ at the initial endpoints of all occurrences of E_r and at the terminal endpoints of all occurrences of E_r^{-1} yields a splitting of ρ into basic paths of height r and paths in G_{r-1} .

Observe that if $H_r = \{E_r\}$ is a polynomially growing stratum, then $f_{\#}^k(E_r) = E_r \cdot u_r \cdot f_{\#}(u_r) \cdot \dots \cdot f_{\#}^{k-1}(u_r)$. Each subpath of the form $f_{\#}^i(u_r)$ is called a *block* of $f_{\#}^k(E_r)$. Since there is no cancellation between successive blocks, it makes sense to refer to the infinite path

$$R_r = u_r \cdot f_{\#}(u_r) \cdot f_{\#}^2(u_r) \cdot \dots \quad (1)$$

as the *eigenray* of E_r .

Remark 1.7 (A note on terminology). The notion of a *polynomially growing stratum* $H_r = \{E_r\}$ first appeared in [BH92]. Polynomially growing strata are called *nonexponentially growing* strata in [BFH00]. Both terms are somewhat misleading because the function $k \mapsto |f_{\#}^k(E_r)|$ may grow exponentially (see Lemma 2.4).

Given an improved train track map $f: G \rightarrow G$, we construct a metric on G . If H_r is an exponentially growing stratum, then its transition matrix M_r has a unique positive left eigenvector v_r (corresponding to λ_r) whose smallest entry equals one [Gan59]. For an edge E_i in H_r , the eigenvector v_r has an entry $l_i > 0$ corresponding to E_i . We choose a metric on G such that E_i is isometric to an interval of length l_i , and such that edges in zero strata or in polynomially growing strata are isometric to an interval of length one. For a path ρ , we denote its length by $\mathcal{L}(\rho)$. Note that if the endpoints of ρ are vertices, then the number of edges in ρ provides a lower bound for $\mathcal{L}(\rho)$. Moreover, if f is an absolute train track map, then f expands the length of legal paths by the factor λ .

Remark 1.8. We merely choose this metric for convenience. All statements here are invariant under bi-Lipschitz maps, but our metric of choice simplifies the presentation of our arguments.

We are now ready to state the main technical result of this paper.

Theorem 1.9. *Let $\phi: F \rightarrow F$ be an automorphism. Then there exists an improved relative train track map representing some positive power of ϕ for which there exists a constant $K \geq 1$ with the following property: For any pair of exponents N, i satisfying $0 \leq i \leq N$, the following two statements hold:*

1. *If σ is a circuit in G , then*

$$\mathcal{L}(f_{\#}^i(\sigma)) \leq K(\mathcal{L}(\sigma) + \mathcal{L}(f_{\#}^N(\sigma))).$$

2. *If ρ is a path in G that starts and ends at vertices, then*

$$\mathcal{L}(f_{\#}^i(\rho)) \leq K(\mathcal{L}(\rho) + \mathcal{L}(f_{\#}^N(\rho))).$$

Given the improved relative train track map $f: G \rightarrow G$, the constant K can be computed.

We will present the proof of Theorem 1.9 in Section 5 and Section 6. Right now, we show how Theorem 0.1 follows from Theorem 1.9.

Proof of Theorem 0.1. Let $\phi: F \rightarrow F$ be an automorphism of a finitely generated free group $F = \langle x_1, \dots, x_n \rangle$. The first part of Theorem 1.9 immediately implies that the first part of Theorem 0.1 holds for some positive power ϕ^k , i.e., there exists some $K' \geq 1$ such that for all $0 \leq i \leq N$ and $w \in F$, we have

$$\|\phi_{\#}^{ik}(w)\| \leq K'(\|w\| + \|\phi_{\#}^{Nk}(w)\|),$$

where we compute lengths with respect to the generators x_1, \dots, x_n .

Let $L = \max\{|\phi(x_i)|\}$. Then, for $0 \leq j < k$, we have

$$L^{-k}\|\phi^{ik+j}(w)\| \leq \|\phi^{ik}(w)\| \leq L^k\|\phi^{ik+j}(w)\|$$

for all $w \in F$. We conclude that for all $0 \leq i \leq N$ and $w \in F$, we have

$$L^{-k}\|\phi_{\#}^i(w)\| \leq K'(\|w\| + L^k\|\phi_{\#}^N(w)\|),$$

so that the first part of Theorem 0.1 holds with $K = L^{2k}K'$.

In order to prove the second assertion, we modify a trick from [BFH97]. Let F' be the free group generated by x_1, \dots, x_n and an additional generator a . We define an automorphism $\psi: F' \rightarrow F'$ by letting $\psi(x_i) = \phi(x_i)$ for all $1 \leq i \leq n$, and $\psi(a) = a$.

By the previous step, the first part of Theorem 0.1 holds for ψ , with some constant $K' \geq 1$. Let w be some word in F . Then, for all $i \geq 0$, $\psi^i(aw)$ is a cyclically reduced word in F' , so that we have $|\phi^i(w)| + 1 = ||\psi^i(aw)||$. We conclude that

$$|\phi^i(w)| + 1 \leq K'(|w| + |\phi^N(w)| + 2),$$

for all $0 \leq i \leq N$. Now the second assertion of Theorem 0.1 holds with $K = 2K'$. \square

2 More on train tracks

Thurston's bounded cancellation lemma is one of the fundamental tools in this paper. We state it in terms of homotopy equivalences of graphs.

Lemma 2.1 (Bounded cancellation lemma [Coo87]). *Let $f: G \rightarrow G$ be a homotopy equivalence. There exists a constant \mathcal{C}_f , depending only on f , with the property that for any tight path ρ in G obtained by concatenating two paths α, β , we have*

$$\mathcal{L}(f_{\#}(\rho)) \geq \mathcal{L}(f_{\#}(\alpha)) + \mathcal{L}(f_{\#}(\beta)) - \mathcal{C}_f.$$

An upper bound for \mathcal{C}_f can easily be read off from the map f [Coo87]. Let $f: G \rightarrow G$ be an improved relative train track map with an exponentially growing stratum H_r with growth rate λ_r . The r -length of a path ρ in G , $\mathcal{L}_r(\rho)$, is the total length of $\rho \cap H_r$.

If β is an r -legal path in G whose r -length satisfies $\lambda_r \mathcal{L}_r(\beta) - 2\mathcal{C}_f > \mathcal{L}_r(\beta)$ and α, γ are paths such that the concatenation $\alpha\beta\gamma$ is an immersion, then the r -length of the segment in $f_{\#}^k(\alpha\beta\gamma)$ corresponding to β (Convention (1.2)) will tend to infinity as k tends to infinity. The *critical length* \mathcal{C}_r of H_r is the infimum of the lengths satisfying the above inequality, i. e.,

$$\mathcal{C}_r = \frac{2\mathcal{C}_f}{\lambda_r - 1}. \quad (2)$$

We now list some additional technical results about improved train track maps. The following lemma is an immediate consequence of [Bri00, Proposition 6.2]. If H_r is an exponentially growing stratum, and ρ is a path of height r , we let $n(\rho)$ denote the number of r -legal segments in ρ .

Lemma 2.2. *Let $f: G \rightarrow G$ be a relative train track map, and let H_r be an exponentially growing stratum. For each $L > 0$, there exists some computable exponent $M > 0$ such that if ρ is a path or circuit in G_r containing at least one full edge in H_r , one of the following three statements holds:*

1. $f_{\#}^M(\rho)$ has an r -legal segment of r -length greater than L .
2. $n(f_{\#}^M(\rho)) < n(\rho)$.
3. ρ can be expressed as a concatenation $\tau_1 \rho' \tau_2$, where τ_1, τ_2 each contain at most one r -illegal turn, the r -length of the r -legal segments of τ_1, τ_2 is at most L , and ρ' splits as a concatenation of pre-Nielsen paths (with one r -illegal turn each) and segments in G_{r-1} . Moreover, $f_{\#}^M(\rho')$ is a concatenation of Nielsen paths of height r and segments in G_{r-1} .

Remark 2.3.

- The statement of Lemma 2.2 in [Bri00] does not explicitly mention the computability of M . The proof, however, only uses counting arguments, from which the constant M can be computed.
- The presence of the subpaths τ_1, τ_2 in Part 3 is an artifact of the fact that ρ need not start or end at fixed points if it is a path. If ρ starts at a fixed point, then τ_1 will be trivial, and if ρ ends at a fixed point, then τ_2 will be trivial.
- The actual statement of [Bri00, Proposition 6.2] does not mention circuits since they were not a concern in the context of [Bri00]. The proof, however, works for circuits as well as paths. If the first two statements of Lemma 2.2 do not hold, then the third statement will hold with τ_1 and τ_2 trivial.

From now on, we assume that $f: G \rightarrow G$ that f is an improved train track map. Throughout the rest of this section, let M be the constant from Lemma 2.2 for some fixed $L > C_r$ (Equation 2).

Let $H_r = \{E_r\}$ be a polynomially growing stratum. We say that H_r is *truly polynomial* if u_r is trivial or, inductively, if u_r is a concatenation of truly polynomial edges and Nielsen paths in exponentially growing strata. Clearly, if E_r is truly polynomial, then the map $k \mapsto |f_{\#}^k(E_r)|$ grows polynomially. We say that a polynomially growing stratum is *fast* if it is not truly polynomial.

The following lemma give us an understanding of the growth of fast polynomial strata.

Lemma 2.4. *There exists an exponent k_0 with the following property: For all fast polynomial strata $H_r = \{E_r\}$ there exists some $s < r$ such that H_s is of exponential growth and $f_{\#}^{k_0}(E_r)$ contains an s -legal subpath of height s whose s -length exceeds C_s .*

In particular, this lemma implies that fast polynomial strata grow exponentially. Given an improved relative train track map, we can find k_0 by successively evaluating $f_{\#}, f_{\#}^2, \dots$ until we see long legal segments in all images of fast polynomial edges.

Proof. We introduce *classes* of fast polynomial edges. Let $H_r = \{E_r\}$ be a fast polynomial edge such that $f(E_r) = E_r u_r$. We say that H_r has *class* 1 if there

exists some $s < r$ such that H_s is an exponentially growing stratum, $u_r \cap H_s$ does not only consist of Nielsen paths and paths of height less than s , and if u_r contains any polynomial edges E_t for some $t > s$, then E_t is truly polynomial. We recursively define a fast polynomial edge E_r to have *class* k if the highest class of edges in u_r is $k - 1$.

If H_r has class 1, then u_r contains a subpath ρ of height s such that $f_{\#}^k(\rho)$ contains a long s -legal segment for some sufficiently large k (Lemma 2.2). If u_r contains any subpaths whose height exceeds s , then by definition those subpaths will grow at most polynomially, so that eventually, the exponential growth of ρ will prevail.

In order to prove the lemma for an edge of class k , $k > 1$, we observe that no edges of class $k - 1$ are cancelled when $f^m(u_r)$ is tightened to $f_{\#}^m(u_r)$. Now the lemma follows by Theorem 1.4, Part 4, and induction. \square

Assume that H_r is an exponentially growing stratum, and let ρ be a path of height r . If H_r does not support a closed Nielsen path, then we let $N(\rho) = n(\rho)$. If H_r supports a closed Nielsen path, then we let $N(\rho)$ equal the number of legal segments in ρ that do not overlap with a Nielsen subpath of ρ .

The following lemma is a generalization of [Bri00, Lemma 6.4].

Lemma 2.5. *Assume that H_r is an exponentially growing stratum. There exist computable constants $\lambda > 1, N_0$ with the following property: If $f_{\#}^M(\rho)$ does not contain a legal segment of length at least L , and if $N(\rho) > N_0$, then*

$$N(f_{\#}^M(\rho)) \leq \lambda^{-1}N(\rho).$$

Regardless of $N(\rho)$, we have

$$N(f_{\#}^M(\rho)) \leq \lambda^{-1}N(\rho) + 1.$$

Proof. If H_r does not support a closed Nielsen path, then the proof of [Bri00, Lemma 6.4] goes through unchanged. We repeat the argument here because the ideas of the proof show up more clearly in this case.

If H_r does not support a closed Nielsen path, then the proof is based on the following observation: If $N(\rho) = 6$ and $f_{\#}(\rho)$ does not contain a long legal segment, then $N(f_{\#}^M(\rho)) < 6$. Suppose otherwise, i.e., $N(f_{\#}^M(\rho)) = N(\rho)$. Then, by Lemma 2.2, $f_{\#}^M(\rho) = \tau_1\gamma\tau_2$, where γ is a concatenation of three indivisible Nielsen paths of height r and paths in G_{r-1} . This is impossible because by Theorem 1.4, Part 3, we can concatenate no more than two indivisible Nielsen paths of height r with paths in G_{r-1} .

Hence, of every six consecutive legal segments in ρ , at least one cancels completely when $f^M(\rho)$ is tightened to $f_{\#}^M(\rho)$. This implies that if $N(\rho) \geq 6$, then $N(f_{\#}^M(\rho)) \leq \frac{10}{11}N(\rho)$. In order to see why this choice of λ works, we just observe that if ρ consists of eleven legal segments and the sixth one cancels in $f_{\#}^M(\rho)$, then there are no six consecutive legal segments that survive in $f_{\#}^M(\rho)$.

This completes the proof of the first inequality, with $\lambda = \frac{11}{10}$ and $N_0 = 5$, if H_r does not support a closed Nielsen path. Regarding the second inequality, we remark that if $N(\rho) \leq N_0$, then $N(f_{\#}^M(\rho)) \leq N(\rho) \leq \lambda^{-1}N(\rho) + 1$.

We now assume that H_r supports a closed indivisible Nielsen path σ . The proof in this case is based on the following consequence of Lemma 2.2. If γ a path of height r , $n(\gamma) = n(f_{\#}^M(\gamma)) = 4$, and $f_{\#}^M(\gamma)$ does not contain a long legal segment, then $f_{\#}^M(\gamma) = \tau_1 \sigma^{\pm 1} \tau_2$, where τ_1 and τ_2 are as in Lemma 2.2. Intuitively, this means that if few legal segments disappear, then many Nielsen paths will appear. Since $N(\rho)$ only counts those legal segments that do not overlap with a Nielsen path, this observation will yield the desired estimate.

First, consider a path γ of height r that does not contain any Nielsen subpaths, i.e., we have $N(\gamma) = n(\gamma)$. If $N(\gamma) \geq 4$, then for every four consecutive legal segments whose images do not cancel completely in $f_{\#}^M(\gamma)$, $f_{\#}^M(\gamma)$ contains at least one Nielsen subpath, so that we have $N(f_{\#}^M(\gamma)) \leq \frac{6}{7}N(\gamma)$, using the same reasoning as above.

We claim that if γ starts and ends at fixed points, then, by Remark 2.3, we have $N(f_{\#}^M(\gamma)) \leq \frac{6}{7}N(\gamma)$ regardless of $N(\gamma)$. To this end, we first argue that if γ starts and ends at fixed points, then $n(f_{\#}^M(\gamma)) < n(\gamma)$. If this were not true, then, by Lemma 2.2 we would have $f_{\#}^M(\gamma) = \sigma^m$ for some $m \in \mathbb{Z}$, which would imply that $\gamma = \sigma^m$ because γ starts and ends at fixed points. This is a contradiction since we assumed that γ does not contain any Nielsen subpaths. Now, if $n(\gamma) = N(\gamma) < 4$, then we conclude that $N(f_{\#}^M(\gamma)) \leq n(f_{\#}^M(\gamma)) < n(\gamma)$. Now $n(\gamma) < 4$ implies that $\frac{6}{7}n(\gamma) \geq n(\gamma) - 1$, which implies that $N(f_{\#}^M(\gamma)) \leq \frac{6}{7}n(\gamma) = \frac{6}{7}N(\gamma)$.

After these preparations, we express ρ as a concatenation

$$\rho = \rho_1 \sigma^{n_1} \rho_2 \sigma^{n_2} \rho_3 \cdots \rho_k \sigma^{n_k} \rho_{k+1},$$

where $n_1, \dots, n_k \in \mathbb{Z}$, and none of the subpaths ρ_i contains a Nielsen subpath.

Note that the subpaths ρ_2, \dots, ρ_k start and end at the base point v of the Nielsen path σ , which is fixed by f . Hence, for $2 \leq i \leq k$, we have $N(f_{\#}^M(\rho_i)) \leq \frac{6}{7}N(\rho_i)$, and we have $N(f_{\#}^M(\rho_1)) \leq \frac{6}{7}N(\rho_1)$ (resp. $N(f_{\#}^M(\rho_{k+1})) \leq \frac{6}{7}N(\rho_{k+1})$) if $N(\rho_1) \geq 4$ (resp. $N(\rho_{k+1}) \geq 4$).

If $N(\rho_1) < 4$ and $N(\rho_{k+1}) < 4$, we have

$$\begin{aligned} N(f_{\#}^M(\rho)) &\leq N(\rho_1) + N(\rho_{k+1}) + \frac{6}{7}(N(\rho) - N(\rho_1) - N(\rho_{k+1})) \\ &\leq 6 + \frac{6}{7}(N(\rho) - 6) \leq \frac{6}{7}(1 + N(\rho)). \end{aligned}$$

Similar estimates yield that $N(f_{\#}^M(\rho)) \leq \frac{6}{7}(1 + N(\rho))$ regardless of $N(\rho_1)$ and $N(\rho_{k+1})$.

If $N(\rho) > 11$, then $\frac{6}{7}(1 + N(\rho)) \leq \frac{13}{14}N(\rho)$, which implies that $N(f_{\#}^M(\rho)) \leq \frac{13}{14}N(\rho)$ if $N(\rho) > 11$, so that the first inequality of the lemma holds with $\lambda = \frac{14}{13}$ and $N_0 = 11$. As for the second inequality, we remark that $N(f_{\#}^M(\rho)) \leq N(\rho)$ and, if $N(\rho) \leq 11$, then $N(\rho) \leq \lambda^{-1}N(\rho) + 1$. \square

The next lemma is a statement about the (absence of) cancellation between eigenrays of polynomially growing strata. It is a stronger version of [BFH00, Sublemma 1, Page 587].




Figure 1: The idea of the proof of Lemma 2.6.

Lemma 2.6. *Let $H_i = \{E_i\}$ and $H_j = \{E_j\}$ be polynomially growing strata. Let S_i (resp. S_j) be an initial segment of $E_i R_i$ (resp. $E_j R_j$, see Equation 1) such that the concatenation $S_i \bar{S}_j$ is a path. If E_i grows faster than linearly and if an entire block of R_j is canceled in $f_{\#}^k(S_i \bar{S}_j)$ for some $k \geq 0$, then no entire block of R_i will be canceled in $f_{\#}^l(S_i \bar{S}_j)$ for any $l \geq 0$.*

Proof. Suppose that at least one block of both S_i and S_j cancels. Then there are paths α , β , and γ such that $f_{\#}^k(u_i) = \beta\gamma$, $f_{\#}^l(u_j) = \alpha\beta$ for some $k, l \geq 0$, and $f_{\#}(\alpha) = \gamma$ (see Figure 1).

In particular, we have

$$R_i = u_i f_{\#}(u_i) \cdots f_{\#}^{k-1}(u_i) \beta f_{\#}(\alpha) f_{\#}(\beta) f_{\#}^2(\alpha) \cdots,$$

and

$$R_j = u_j f_{\#}(u_j) \cdots f_{\#}^{l-1}(u_j) \alpha \beta f_{\#}(\alpha) f_{\#}(\beta) f_{\#}^2(\alpha) \cdots.$$

In particular, the path $\rho = E_i R_i^{k-1} \bar{\alpha} \bar{R}_j^{l-1} \bar{E}_j$ does not split. By Theorem 1.4, ρ is a exceptional path, and both E_i and E_j grow linearly. \square

3 Terminology and examples

In this section, we discuss some examples that illustrate some of the main issues that we need to address in the proof of Theorem 1.9. Although we are not primarily concerned with free-by-cyclic groups in this article, the language of free-by-cyclic groups will streamline the exposition.

Given a free group $F_n = \langle x_1, \dots, x_n \rangle$ and an automorphism ϕ of F_n , the *mapping torus* of ϕ is the free-by-cyclic group

$$M_{\phi} = \langle x_1, \dots, x_n, t \mid t^{-1} x_i t \phi(x_i^{-1}) \rangle.$$

The letter t is called the *stable letter* of M_{ϕ} .

A reduced word w in the generators of M_{ϕ} is a *hallway* if w represents the trivial element of M_{ϕ} and if w can be expressed as $w = w_1 w_2$ such that w_1 only contains negative powers of t and w_2 only contains positive powers of t [BF92]. Hallways of the form $t^{-k} x t^k \phi^k(x^{-1})$, for $x \in F_n$, are said to be *smooth*.




Figure 2: A hallway.

Any hallway w can be expressed as

$$w = t^{-1}u_{k-1}t^{-1}u_{k-2}t^{-1} \cdots t^{-1}u_1t^{-1}w_0tv_1tv_2t \cdots tv_{k-1}tw_k^{-1},$$

where $w_0, w_k, u_1, \dots, u_{k-1}, v_1, \dots, v_{k-1}$ are elements of F_n . The words u_i and v_i may be empty. In fact, a hallway is smooth if and only if all the u_i and v_i are trivial. For $1 \leq i < k$, we define w_i to be the word obtained by tightening $u_i\phi(w_{i-1})v_i$. Since w represents the identity, we have $w_k = \phi(w_{k-1})$. We call w_i the i -th *slice* of w . The number k is the *duration* $\mathcal{D}(w)$ of the hallway. Figure 2 illustrates these notions.

We say that the instances of letters of F_n that occur in the spelling of w are *visible*. Theorem 0.1 states that if w is a smooth hallway, then the length of each w_i is bounded by a constant multiple of the number of visible edges in w .

The following examples illustrate the main issues that arise in the proof. For the remainder of this section, let $F_6 = \langle a, b, c, d, x, y \rangle$, and define ϕ by letting

$$\begin{aligned} a &\mapsto a \\ b &\mapsto ba \\ c &\mapsto caa \\ d &\mapsto dc \\ x &\mapsto y \\ y &\mapsto xcy. \end{aligned}$$

This automorphism admits the stratification $H_1 = \{a\}$, $H_2 = \{b\}$, $H_3 = \{c\}$, $H_4 = \{d\}$, and $H_5 = \{x, y\}$. The restriction of f to the filtration element $G_3 = H_1 \cup H_2 \cup H_3$ grows linearly, the restriction to G_4 grows quadratically, and the stratum H_5 is of exponential growth.




Figure 3: Illustration of Example 3.1.




Figure 4: Illustration of Example 3.2.

The first example illustrates the behavior of smooth hallways in linearly growing filtration elements.

Example 3.1. Let w_0 be a word from the list $a^m, ba^m b^{-1}, ca^m c^{-1}$, for some integer m . Then $\phi(w_0) = w_0$, so that the length of any slice of the hallway $t^{-k}w_0 t^k w_0^{-1}$ is the same as the length of w_0 . Now, let w_0 be a word from the list $ba^m, ca^m, ca^m b^{-1}$. If $m \geq 0$, then $|\phi^{k+1}(w_0)| = |\phi^k(w_0)| + 1$ for any $k \geq 0$. If $m < 0$, then $|\phi^{k+1}(w_0)| = |\phi^k(w_0)| - 1$ for $0 \leq k < -m$ (Figure 3). Hence, the length of each slice of the hallway $t^{-k}w_0 t^k \phi^k(w_0^{-1})$ is bounded by the number of visible letters.

If w_0 is an arbitrary word in $\langle a, b, c \rangle$, then, by Remark 1.6, it splits as a concatenation of words from the above lists and their inverses, which implies that the lengths of slices of smooth hallways is bounded by the number of visible letters, so that Theorem 0.1 holds with $K = 1$.

The next example shows that hallways that are not smooth may have slices whose length is not bounded in terms of a constant multiple of the number of visible edges.

Example 3.2. Let $w = t^{-k} c t^{-k} b^{-1} t^{2k} b c^{-1}$. For $i < k$, we have $w_i = a^{-i} b^{-1}$, and for $k \leq i \leq 2k$, we have $w_i = c a^{2k-i} b^{-1}$ (Figure 4). In particular, there is a




Figure 5: Illustration of Example 3.3.

slice of length $k + 2$ although there are only four visible edges in w . Informally, one might say that hallways of this form bulge in the middle. A similar bulge occurs for hallways of the form $w = t^{-k}b^{-1}t^{-k}bt^kb^{-1}t^kb$.

The next example shows that we need to control the size of such bulges when proving Theorem 1.9.

Example 3.3. First, note that for $k \geq 1$, the last letter in the words $f^{-k}(xc)$ is always one of x, y, x^{-1}, y^{-1} , so that words of the form $w_0 = \phi^{-k}(xc)b^{-1}$ are reduced, and we have $\phi^k(w_0) = xca^kb^{-1}$ and $\phi^{2k}(w_0) = \phi^k(x)b^{-1}$.

Hence, the smooth hallway $w = t^{-2k}w_0t^{2k}\phi^{2k}(w_0^{-1})$ contains a bulge like the first one in the previous example (Figure 5). The presence of this bulge does not contradict Theorem 0.1 because w contains a large number of visible instances of the letters x and y . This example shows that we cannot consider the strata separately when proving Theorem 1.9.

Example 3.4. If we let $w_0 = \phi^{-k}(dc)b^{-1}$, then the smooth hallway $w = t^{-2k}w_0t^{2k}\phi^{2k}(w_0^{-1})$ contains a bulge like in Example 3.3. This does not contradict Theorem 0.1 as w contains a large number of visible instances of the letter c .

Our final example illustrates a subtlety regarding linearly growing strata.

Example 3.5. Let $F_4 = \langle a, b, c, d \rangle$ and define ψ by letting

$$\begin{aligned} a &\mapsto a \\ b &\mapsto ba \\ c &\mapsto ca \\ d &\mapsto dcb^{-1}. \end{aligned}$$

The map ψ is a linearly growing automorphism, so in particular the letter d is of linear growth, although the image of d contains letters of linear growth other than d itself.

Letters of linear growth may thus behave in two different ways; they may contribute to the growth of images under successive applications of ψ , or they may remain inert as parts of a fixed word. In the proof of Theorem 1.9, we will need to distinguish letters of linear growth according to their role.

Example 3.6. Let $F_3 = \langle a, x, y \rangle$ and define ϕ by letting

$$\begin{aligned} a &\mapsto axyx^{-1}y^{-1} \\ x &\mapsto y^{-1} \\ y &\mapsto yx. \end{aligned}$$

The stratum $\{x, y\}$ grows exponentially, and we have $\psi(xyx^{-1}y^{-1}) = axyx^{-1}y^{-1}$. This means that a grows linearly although it maps across an exponentially growing stratum. This is another phenomenon that we need to consider when analyzing strata of linear growth.

The notion of hallways naturally extends to mapping tori of homotopy equivalences of finite graphs. Specifically, a hallway ρ in the mapping torus of $f: G \rightarrow G$ is a sequence of paths of the form

$$\rho = (\mu_{k-1}, \mu_{k-2}, \dots, \mu_1, \rho_0, \nu_1, \nu_2, \dots, \nu_{k-1}, \rho_k),$$

where $\rho_0, \rho_k, \mu_1, \dots, \mu_{k-1}, \nu_1, \dots, \nu_{k-1}$ are paths in G , satisfying $f(\tau(\rho_0)) = \iota(\nu_1)$, $f(\tau(\nu_i)) = \iota(\nu_{i+1})$, $f(\tau(\nu_{k-1})) = \tau(\rho_k)$, $f(\iota(\rho_0)) = \tau(\mu_1)$, $f(\iota(\mu_i)) = \tau(\mu_{i+1})$, and $f(\iota(\mu_{k-1})) = \iota(\rho_k)$, where $\iota(\cdot)$ is the initial point of a path, and $\tau(\cdot)$ is the terminal point.

The paths μ_i and ν_i are called *notches*. Some or all of the notches may be trivial. For $1 \leq i < k$, we define ρ_i to be the path obtained by tightening $\mu_i f(\rho_{i-1}) \nu_i$. Since ρ is a closed path, we have $\rho_k = f_{\#}(\rho_{k-1})$. As before, we call ρ_i the i -th *slice* of ρ , and the number k is the duration $\mathcal{D}(\rho)$.

The visible length of ρ is

$$\mathcal{V}(\rho) = \mathcal{L}(\rho_0) + \mathcal{L}(\rho_k) + \sum_{i=1}^{k-1} (\mathcal{L}(\mu_i) + \mathcal{L}(\nu_i)).$$

Finally, we introduce *quasi-smooth* hallways: Given some $C \geq 0$, we say that a hallway ρ is C -quasi-smooth if the length of all the notches is bounded by C .

4 Strata of superlinear growth

Throughout this section, let $f: G \rightarrow G$ be an improved relative train track map.

In order to track images of edges through the slices of a hallway ρ , we assign a *marking* to each edge. This assignment will, in general, involve arbitrary choices, but our arguments will not be affected by these choices.

Definition 4.1. We begin by marking all edges in the initial slice ρ_0 and in all notches μ_i, ν_i with their height. Assume inductively that all edges in a slice ρ_{i-1} have been marked, and let E be an edge of height r in ρ_{i-1} , with marking s . Now, consider an edge E' in $f(E)$. If the height of E' is r , or if H_s is a zero stratum, then we keep the marking s . If the height of E' is less than r , then we mark E' by r . This gives us a marking for all edges in $\mu_i f(\rho_{i-1}) \nu_i$.

Note that, as we tighten $\mu_i f(\rho_{i-1}) \nu_i$ to obtain ρ_i , different choices in cancellation (Remark 1.1) may give rise to different possible markings, but this will not be a problem.

We say that an edge E is marked by a linear/polynomial/exponential stratum if its marking is s and H_s is linear/polynomial/exponential.

The following proposition goes a long way toward proving Theorem 1.9. In fact, if f has no edges of linear growth, then it immediately implies Theorem 1.9.

Proposition 4.2. *There exists some constant $K \geq 1$ such that for every hallway ρ and every slice ρ_i of ρ , the number of edges in ρ_i that are not marked by strata of linear growth is bounded by $KV(\rho)$.*

Given the improved relative train track map $f: G \rightarrow G$, the constant K can be computed.

In order to streamline the exposition, we will not always make the choice of K explicit. However, it will turn out that K can be chosen to be the product of numbers that can easily be read off from the train track map.

The intuition of the proof is that once significant growth occurs, it will be due to the presence of long legal subpaths in exponentially growing strata or long subsegments of eigenrays of polynomially growing strata that grow faster than linearly. Lemma 2.1 and Lemma 2.6 imply that there is hardly any cancellation between such subpaths and their surroundings, so that any significant growth that occurs in a slice will eventually be accounted for by visible edges.

The following definition will help us understand cancellation in hallways. For every stratum H_r , we define a number $h(H_r)$ in the following way:

- If H_r is a constant stratum, then $h(H_r) = 0$.
- If H_r is a nonconstant polynomially growing stratum, i.e., $H_r = \{E_r\}$ and $f(E_r) = E_r u_r$, then $h(H_r)$ is the height of u_r .
- If H_r is of exponential growth and H_{r-1} is not a zero stratum, then $h(H_r)$ is the height of $f(H_r) \cap G_{r-1}$, unless this intersection does not contain any edges, in which case we let $h(H_r) = \infty$.
- If H_r is of exponential growth and H_{r-1} is a zero stratum, then $h(H_r)$ is the height of $f(H_r \cup H_{r-1}) \cap G_{r-2}$. We also let $h(H_{r-1}) = h(H_r)$.

Essentially, $h(H_r)$ is the index of the highest stratum crossed by the image of H_r , other than H_r itself. We may permute the strata of G (while preserving the improved train track properties) such that $h(H_r) > h(H_s)$ implies $r > s$.

Given a stratum H_s , we say that the set $S(H_s) = \{H_r | h(H_r) = s\}$ is the *league* of H_s , the motivation being that they, in a sense, “play at the same level.” If $h(H_r) = \infty$, then H_r does not belong to any league.

Proof of Proposition 4.2. First of all, we note that if a slice ρ_i has a subpath in a zero stratum H_r , then this subpath is of uniformly bounded length, and it is surrounded by edges in higher strata (Theorem 1.4, Part 1), so that we have a linear estimate of the number of edges in H_r in ρ_i in terms of the number of edges in higher strata.

Let q be the largest (finite) number for which the league $S(q)$ is nonempty. Fix some stratum H_r for $r > q$. We want to find a linear bound on the number of edges in $\rho_i \cap H_r$ in terms of visible edges. By definition of $S(q)$ and choice of r , edges in $\rho_i \cap H_r$ never cancel with edges from other strata or their images.

If $H_r = \{E_r\}$ is of polynomial growth, then any occurrence of E_r in ρ_i is the image of a visible copy of E_r , and ρ_i contains at most one copy of E_r for each visible copy of E_r . Hence, the number of edges in $\rho_i \cap H_r$ is bounded by the number of visible edges.

Now, assume that H_r is an exponentially growing stratum. A slice ρ_i decomposes into r -legal subpaths with r -illegal turns in between. By Lemma 2.1, a subpath whose r -length is greater than C_r (Equation 2) will eventually be accounted for by visible edges since it will not be shortened by cancellation within slices.

Edges in H_r whose r -distance from an illegal turn is less than $\frac{C_r}{2}$ may cancel eventually, and ρ_i contains at most C_r of them per r -illegal turn, so that we only need to find a bound of the number of r -illegal turns in terms of the number of visible edges. Since the improved train track map f does not create any r -illegal turns, any r -illegal turn in ρ_i can be traced back to a visible illegal turn in ρ (or an illegal turn created by appending a notch to the image of a slice). This implies that the number of r -illegal turns in ρ_i is bounded by the number of visible edges in ρ .

Summing up, we have bounded the number of edges in $\rho_i \cap (H_{q+1} \cup H_{q+2} \cup \dots)$ by a multiple of the number of visible edges. This establishes the base case of the proof.

We now assume inductively that the number of edges in $S(p) \cup S(p+1) \cup \dots$ has been bounded as a constant multiple of $\mathcal{V}(\rho)$. We need to find a bound on the number of edges in $\rho_i \cap H_p$.

We first assume that $H_p = \{E_p\}$ is of polynomial growth. By definition of $S(p)$, an edge in $\rho_i \cap H_p$ has one of four possible markings:

- Its marking may be p , indicating that it is the image of a visible edge, or
- it may be marked by an exponentially growing stratum in $S(q)$, for some $q \geq p$, or
- it may be marked by a superlinear polynomially growing stratum in $S(q)$, $q \geq p$, or
- it may be marked by a stratum of linear growth.

We are not concerned with edges of the fourth kind.

As before, the number of edges of the first kind in $\rho_i \cap H_p$ is bounded by the number of visible edges. Let C be the largest number of copies of E_p that occur in the image of a single edge in an exponentially growing stratum H_s , for $s > p$. Then the number of edges of the second kind in $\rho_i \cap H_p$ is bounded by C times the number of exponentially growing edges in $\rho_{i-1} \cap (S(p) \cap S(p+1) \cap \dots)$, which in turn is bounded by a multiple of the number of visible edges.

We have no immediate bound on the number of edges of the third kind. As we trace the image of such an edge through subsequent slices, one of three possible events will occur:

- Either, it eventually maps to a visible edge, or
- it cancels with an edge of the first or second kind, or
- it cancels with an edge in the image of a polynomially growing (possibly linearly growing) edge in $S(p)$.

Note that these events may depend on choices in tightening (Remark 1.1), but once again our estimates will not be affected by these choices.

The number of edges for which one of the first two events occurs is clearly bounded by a multiple of the number of visible edges. We only need to find a bound on the number of edges in an eigenray that eventually cancel with edges in another eigenray.

Lemma 2.6 implies that there is a uniform bound on the number of edges in H_p that cancel when two rays meet, so that we only need to find a bound on the number of meetings between two rays. Clearly, any two rays meet at most once.

If an eigenray cancels with segments from more than one other ray (this is conceivable since a slice may be of the form $\rho_i = E_r S_1 S_2$, where E_r is a polynomially growing edge in $S(p)$ and S_1, S_2 are short segments from rays of edges in $S(p)$ such that the ray of E_r successively cancels with S_1 and S_2), then all except possibly one of these segments cancel completely, so that they are no longer available for subsequent cancellation. This implies that the number of meetings of rays is bounded by two times the number of pieces of rays available for cancellation, which in turn is bounded by the number of visible edges.

This completes our estimate of the number of edges in $\rho_i \cap H_p$ when H_p is of polynomial growth. We now assume that H_p is of exponential growth.

The number of subpaths of height p of ρ_i is bounded by the number of edges of height greater than p in ρ_i plus one. The contribution of p -legal subpaths of p -length less than or equal to C_p is bounded by C_p times the number of subpaths of height p , so that we do not need to consider them here. Any p -legal subpaths of length greater than C_p will eventually show up in the visible part of ρ , so that we do not need to consider them, either. The remaining edges in $\rho_i \cap H_p$ are at p -distance less than $\frac{C_p}{2}$ from a p -illegal turn. Hence, we only need to find a bound on the number of p -illegal turns in ρ_i .

As before, we trace illegal turns in ρ_i back to their origin:

- An illegal turn may be the image of a visible illegal turn (this case includes illegal turns created by appending notches μ_i, ν_i to the image $f_{\#}(\rho_{i-1})$ of a slice), or
- it may come from a illegal turn in the image of an exponentially growing edge in $S(p)$, or
- it may be contained in the ray of a polynomially growing edge in $S(p)$, or
- it may be contained in a Nielsen path marked by a linear stratum (Example 3.6).

We are not concerned with illegal turns of the fourth type.

The same arguments that we used for polynomially growing H_p yield that the number of illegal turns of the first and second kind is bounded by a multiple of the number of visible edges.

Now, let C be the maximum of the number of illegal turns in the images of polynomially growing edges in $S(p)$. Lemma 2.4 yields an exponent k_0 such that for polynomially growing edge E_r in $S(p)$, $f_{\#}^{k_0}(u_r)$ contains a long legal segment. This means, in particular, that if ρ contains a block $f_{\#}^k(u_r)$, $k \geq k_0$, then this block contains no more than C illegal turns per long legal segment. Since long legal segments eventually show up as visible edges, the number of illegal turns in such blocks is bounded by $C\mathcal{V}(\rho)$.

The remaining illegal turns are contained in initial subpaths of rays that contain no more than the first $C + 1$ blocks, i.e., there are at most $C(C + 1)$ illegal turns of this kind per ray. Since we already know that the number of rays is bounded in terms of the number of visible edges, we are done in this case.

We have now obtained the desired estimate for edges in ρ_i of height p and higher. In particular, this includes all strata in $S(p - 1)$, which completes the inductive step. \square

5 Polynomially growing automorphisms

In this section, we establish Theorem 1.9 in the case of polynomially growing automorphisms. Specifically, we find estimates for the contribution of linearly growing edges that we ignored in Proposition 4.2. As usual, let $f: G \rightarrow G$ be an improved relative train track map. Since f is of polynomial growth, every stratum H_r contains only one edge E_r , and we have $f(E_r) = E_r \cdot u_r$, where u_r is some closed path in G_{r-1} . Note that all vertices of G are fixed.

We first record an obvious lemma.

Lemma 5.1. *Let μ_1, μ_2 be Nielsen paths in G , and let ν be some path in G .*

- *If μ_1 and μ_2 can be concatenated, then the path obtained from $\mu_1\mu_2$ by tightening relative endpoints is also a Nielsen path.*

- If μ_1 and ν can be concatenated, let γ be the path obtained by tightening $\mu_1\nu$, and let $\Delta = L(\gamma) - L(\nu)$. Then, for all $k \geq 0$, we have

$$L(f_{\#}^k(\gamma)) = L(f_{\#}^k(\nu)) + \Delta,$$

and

$$-L(\mu_1) \leq \Delta \leq L(\mu_1).$$

□

We now establish Theorem 1.9 for automorphisms of linear growth. This lemma will provide the base case of our inductive proof of Theorem 1.9.

Lemma 5.2. *Assume that $f: G \rightarrow G$ is of linear growth. If ρ is a smooth hallway, and if ρ_0 starts and ends at vertices, then the lengths of slices of ρ are bounded by $V(\rho)$, i.e., Theorem 1.9 holds with $K = 1$.*

Proof. The proof proceeds by induction up through the strata of G . The bottom stratum H_1 is constant, so that the lemma trivially holds for the restriction of f to H_1 . We now assume that H_r is a linearly growing stratum, and that the lemma holds for the restriction of f to G_{p-1} .

Consider the initial slice ρ_0 . Remark 1.6 yields a splitting of ρ_0 into basic paths of height p and paths in G_{p-1} . The splitting of ρ_0 induces a decomposition of ρ into smooth hallways, so that it suffices to prove the claim for hallways whose initial slice is a basic path of height p or a path in G_{p-1} .

By induction, we only need to prove the claim if ρ_0 is a basic path of height p . If the basic path ρ_0 is, in fact, an exceptional path, then the reasoning of Example 3.1 proves our claim, so that we may assume that ρ_0 is not an exceptional path.

Assume that ρ_0 is a basic path of the form $E_p\gamma$. Then, by Theorem 1.4, Part 4, there exists some smallest exponent $m \geq 0$ for which $f_{\#}^{m+1}(E_p\gamma)$ splits as $E_p \cdot \gamma'$. Using Remark 1.6 once more, we conclude that $E_p\gamma$ can be expressed as $E_p u_p^{-m} \nu$.

If $D(\rho) \leq m$, then ρ_0 k -splits as $E_p u_p^{-m} \cdot \nu$. We can consider the subpaths $E_p u_p^{-m}$ and ν separately, so that we are done in this case.

Now assume that $D(\rho) > m$. For $0 \leq i \leq m$, we have $\rho_i = E_p u_p^{i-m} f_{\#}^i(\nu)$ and $L(\rho_i) = 1 + (m-i)L(u_p) + L(f_{\#}^i(\nu))$. For $m+1 \leq i \leq D(\rho)$, we have $\rho_i = E_p u_p^{i-(m+1)} f_{\#}^i(u_p \nu)$ and $L(\rho_i) = 1 + (i-m-1)L(u_p) + L(f_{\#}^i(\nu)) + \Delta$, where Δ is defined as in Lemma 5.1.

We have $V(\rho) = 2 + (D(\rho) - 1)L(u_p) + L(\nu) + L(f_{\#}^{D(\rho)}(\nu)) + \Delta$. By induction, we have $L(f_{\#}^i(\nu)) \leq L(\nu) + L(f_{\#}^{D(\rho)}(\nu))$ for all $0 \leq i \leq D(\rho)$. This immediately implies that $L(\rho_i) \leq V(\rho)$ for all $0 \leq i \leq D(\rho)$.

If $\rho_0 = E_p \gamma E_p^{-1}$, we essentially repeat the same argument. Once more, we can write $\rho_0 = E_p u_p^{-m} \nu$, and in order to use the previous argument, we only need to know that the lemma holds for ν . This, however, follows from the previous step, so that we are done. □




Figure 6: Cutting and the sawtooth construction.

We now find estimates on the number of edges emitted by linearly growing edges, the quantity we ignored in Proposition 4.2. The idea is to take a hallway and decompose it into smaller and smaller pieces until all remaining pieces only involve linearly growing edges and their rays. Simple counting arguments will give us bounds on the number of the remaining pieces as well as the lengths of their slices.

Let ρ be a hallway, and assume that there is a visible edge E_r that does not cancel within ρ , i.e., we can trace its image through the slices of ρ until it reappears as another visible edge. Then ρ can be expressed as $\rho = \alpha E_r \beta E_r^{-1}$, and we define two new hallways ρ', ρ'' by tightening $t^{-k} E_r \beta E_r^{-1}$ and αt^k . We say that ρ' and ρ'' are obtained from ρ by *cutting* along the trajectory of E_r (Figure 6). The exponent k is the *length* of the cut. We say that a hallway ρ is *indecomposable* if it does not admit any cuts of length $D(\rho)$.

Now we obtain a new hallway σ from ρ' by repeatedly replacing subwords of the form $t^{-1} E_r$ by $f(E_r) t^{-1}$ and tightening (Figure 6). We refer to this operation as the *sawtooth construction* along the trajectory of E_r .

If \mathcal{M} is a collection of hallways, we let

$$V(\mathcal{M}) = \sum_{\sigma \in \mathcal{M}} V(\sigma).$$

The following lemma lists some basic properties of our two operations. We say that an edge is of *degree* d if $f_{\#}^k(E)$ grows polynomially of degree d .

Lemma 5.3. *Fix some $C \geq \max\{L(u_r)\}$. Let ρ be a C -quasi-smooth hallway in G . Choose $d > 1$ such that the fastest growing edge crossed by ρ grows polynomially of degree d .*

Obtain a collection \mathcal{M} of hallways by cutting along all trajectories of edges E in ρ of degree d . Let \mathcal{M}_1 be the collection of smooth elements of \mathcal{M} , and let \mathcal{M}_2 consist of hallways obtained by performing the sawtooth construction along all trajectories of E of degree d in those elements of \mathcal{M} that are not smooth. Then

1. *The duration of all elements of \mathcal{M}_1 and \mathcal{M}_2 is at most $D(\rho)$.*

2. None of the elements of \mathcal{M}_2 crosses edges of degree d , i.e., they only cross edges of degree at most $d - 1$.
3. All elements of \mathcal{M}_2 are $2C$ -quasi-smooth.
4. The number of elements of \mathcal{M}_2 is bounded by $2CD(\rho)$.
5. We have

$$V(\mathcal{M}_1) + V(\mathcal{M}_2) \leq V(\rho) + (2CD(\rho))^2.$$

Proof. The first four properties follow immediately from definitions. In order to prove the fifth property, we just remark that each element of \mathcal{M}_2 has at most $2CD(\rho)$ visible edges that do not appear in ρ itself. Since \mathcal{M}_2 contains at most $2CD(\rho)$ hallways, the estimate follows. \square

Lemma 5.4. *There exists a (computable) constant C with the following property:*

Let γ be a path of height r , starting and ending at vertices, and assume that E_r is of degree $d > 1$. Then, for all $k \geq 0$,

$$L(\gamma) + L(f_{\#}^k(\gamma)) \geq Ck^d.$$

Proof. It suffices to prove the lemma if either $\gamma = E_r\gamma'$, or $\gamma = E_r\gamma'E_s^{-1}$, where γ' only involves edges of degree less than d , and E_s is of degree d .

In the first case, the claim is obvious. In the second case, we remark that Lemma 2.6 guarantees that there is hardly any cancellation between the rays of E_r and E_s , so that the lemma follows. \square

The following proposition implies the second part of Theorem 1.9 in the case of polynomially growing automorphisms. In particular, it provides bounds on the number of edges emitted by linearly growing edges. This is the quantity that we ignored in Proposition 4.2.

Proposition 5.5. *Assume that f represents an automorphism that grows polynomially of degree q . Fix some $C \geq \max\{L(u_r)\}$. There exist computable constants $K_1 \leq K_2 \leq \dots \leq K_q$ and $K'_1(C), \dots, K'_q(C)$ such that*

1. *If ρ is a smooth hallway whose fastest growing edge is of degree d , and if ρ_0 starts and ends at vertices, then*

$$L(\rho_i) \leq K_d V(\rho)$$

for all slices ρ_i of ρ .

2. *If ρ is a C -quasi-smooth hallway whose fastest growing edge is of degree d , then in every slice ρ_i , the number of edges emitted by linearly growing edges is bounded by*

$$K_d V(\rho) + K'_d(C) D^{d+1}(\rho),$$

so that we have

$$L(\rho_i) \leq (K + K_d)V(\rho) + K'_d(C)D^{d+1}(\rho),$$

where K is the constant from Proposition 4.2.

Proof. We prove the proposition by induction on d . For $d = 1$, the first part holds with $K_1 = 1$ because of Lemma 5.2. Now, assume that ρ is a C -quasi-smooth hallway whose fastest growing edge grows of degree $d = 1$. Obtain a collection \mathcal{M} of hallways by cutting ρ along the trajectories of all linearly growing edges that do not cancel within ρ . If σ is a smooth element of \mathcal{M} , then the first part implies that the number of edges in each σ_i emitted by linearly growing edges is bounded by $V(\sigma)$.

If σ is not smooth, then in every slice σ_i , the number of edges emitted by linearly growing edges is bounded by $V(\sigma) + 2CD(\rho)$ (It is helpful to keep Example 3.2 in mind). Lemma 5.3 yields that \mathcal{M} contains no more than $2CD(\rho)$ pieces that are not smooth. Summing up, we conclude that every slice of ρ contains at most $V(\rho) + (2CD(\rho))^2$ edges emitted by linearly growing edges, so that the second statement follows with $K_1 = 1$ and $K'_1(C) = 4C^2$.

Now, let K be the constant from Proposition 4.2, and assume inductively that the proposition holds for some $d \geq 1$. We want to find some K_{d+1} such that for all hallways ρ whose fastest growing edge is of degree $d + 1$, we have

$$L(\rho_i) \leq K_{d+1}V(\rho).$$

for all slices ρ_i . It suffices to prove this with the assumption that ρ is indecomposable. Then we can perform the sawtooth construction along all trajectories of edges of degree $d + 1$. Since ρ is indecomposable, we obtain one C -quasi-smooth piece σ that only crosses edges of degree d or lower, so that by induction, we conclude that the number of edges in σ_i that were emitted by linearly growing edges is bounded by

$$K_dV(\sigma) + K'_d(C)D^{d+1}(\rho).$$

We conclude that

$$\begin{aligned} L(\rho_i) &\leq KV(\rho) + K_dV(\sigma) + K'_d(C)D^{d+1}(\rho) \\ &\leq (K + K_d)V(\rho) + (2C + K'_d(C))D^{d+1}(\rho). \end{aligned}$$

Using Lemma 5.4, we can find some constant M such that

$$MV(\rho) \geq (2C + K'_d(C))D^{d+1}(\rho)$$

for all indecomposable hallways ρ involving edges of degree $d + 1$. We conclude that the first statement of the proposition holds with $K_{d+1} = K + K_d + M$.

We now prove the second assertion. Let ρ be a C -quasi-smooth hallway. We obtain two collections $\mathcal{M}_1, \mathcal{M}_2$ of hallways by performing cutting and sawtooth operations as in Lemma 5.3.

The elements of \mathcal{M}_1 are smooth hallways, so that for any $\sigma \in \mathcal{M}_1$, the previous step yields

$$L(\sigma_i) \leq K_{d+1}V(\sigma).$$

If σ is an element of \mathcal{M}_2 , then it is a $2C$ -quasi-smooth hallway, and induction yields that in every slice of σ , the number of edges emitted by linearly growing edges is bounded by

$$K_dV(\sigma) + K'_d(2C)D^{d+1}(\sigma).$$

Summing over all elements of \mathcal{M}_1 and \mathcal{M}_2 , we conclude that every slice ρ_i of ρ contains at most

$$\begin{aligned} K_{d+1}V(\mathcal{M}_1) + K_dV(\mathcal{M}_2) + 2CD(\rho) \cdot K'_d(2C)D^{d+1}(\rho) \\ \leq K_{d+1}V(\rho) + 4C^2K_dD^2(\rho) + 2CK'_d(2C)D^{d+2}(\rho) \end{aligned}$$

edges emitted by linearly growing edges, so that the second statement of the proposition holds with

$$K'_{d+1}(C) = 4C^2K_d + 2CK'_d(2C).$$

□

Remark 5.6. The estimates of Proposition 5.5 are rather crude; lots of edges are counted several times rather than just once. I opted to present the most straightforward estimates rather than tightest ones.

6 Proof of the main result

We now extend the techniques and results of Proposition 5 to arbitrary automorphisms. The presence of exponentially growing strata will turn out to be a mixed blessing. On the one hand, they make for rather simple counting arguments as polynomial contributions as in Proposition 5.5 are easily dwarfed by exponential growth. On the other hand, we will need to consider more complicated decompositions of hallways.

As usual, let $f: G \rightarrow G$ be an improved relative train track map. Any statements regarding the computability of constants assume that we are given such a map. After permuting the strata as necessary, we may assume that if H_r and H_s are truly polynomial strata and $r > s$, then the degree of H_r is at least as large as that of H_s . Throughout this section, let K be the constant from Proposition 4.2.

If H_r is an exponentially growing stratum, then we fix some $L > C_r$, and we replace f by f^M , where M is the exponent from Lemma 2.2 for this choice of L . After replacing f by a power yet again if necessary, we may assume that the image of each edge in H_r contains at least L edges in H_r . If H_r supports a closed Nielsen path τ , then the initial and terminal edges of τ are partial edges in H_r , and we may assume that the image of each of them also contains at least L edges in H_r . We say that a legal path of height r is *long* if it contains at least L edges in H_r .

We first record an exponential version of Lemma 5.4.

Lemma 6.1. *Let H_r be an exponentially growing stratum or a fast polynomial stratum. Then there exists a computable constant $\lambda > 1$ such that if σ is a circuit in G_r or a path starting and ending at fixed vertices, then either σ is a concatenation of Nielsen paths of height r and subpaths in G_{r-1} , or we have*

$$\mathcal{L}(\sigma) + \mathcal{L}(f_{\#}^k(\sigma)) \geq \lambda^k$$

for all $k \geq 0$.

Proof. If H_r is an exponentially growing stratum, we need to distinguish two cases: First, assume that for some $i \geq 0$, $f_{\#}^i(\sigma)$ is a concatenation of Nielsen paths and subpaths in G_{r-1} . Since σ starts and ends at fixed vertices, we conclude that σ itself is a concatenation of Nielsen paths and subpaths in G_{r-1} , so that there is nothing to show in this case.

Let λ_{-}, N_0 be the constants from Lemma 2.5, and assume that for all $i \geq 0$, $f_{\#}^i(\sigma)$ is not a concatenation of Nielsen paths and subpaths in G_{r-1} . Let i_0 be the smallest index for which $f_{\#}^{i_0}(\sigma)$ contains a long legal segment. Then, using Lemma 2.5 and Lemma 2.2, we see that $\mathcal{L}(\sigma) \geq \lambda_{-}^{i_0}$. Moreover, we have $\mathcal{L}(f_{\#}^k(\sigma)) \geq \lambda_r^{k-i_0}$.

If we let $\lambda = \sqrt{\min\{\lambda_{-}, \lambda_r\}}$, then we have $\lambda_{-}^{i_0} + \lambda_r^{k-i_0} \geq \lambda^k$. Hence, we have $\mathcal{L}(\sigma) + \mathcal{L}(f_{\#}^k(\sigma)) \geq \lambda_{-}^{i_0} + \lambda_r^{k-i_0} \geq \lambda^k$.

If $H_r = \{E_r\}$ is a fast polynomial stratum, then we argue similarly, using Lemma 2.4 and Theorem 1.4, Part 4. \square

If H_r is an exponentially growing stratum, we let T_r equal the length of the longest path in $f(H_r) \cap G_{r-1}$. We fix another constant $S_r > 0$ with the following property: Let γ be a path in G_{r-1} . If $\mathcal{L}(\gamma) \geq S_r$, then $\mathcal{L}(f_{\#}(\gamma)) > 3T_r$ and $\mathcal{L}(f_{\#}^2(\gamma)) > 3T_r$, and if $\mathcal{L}(\gamma) \leq T_r$, then $\mathcal{L}(f_{\#}(\gamma)) < S_r$ and $\mathcal{L}(f_{\#}^2(\gamma)) < S_r$. We can easily compute a suitable value S_r given the train track map f . We say that a path γ in G_{r-1} is *r-significant* if $\mathcal{L}(\gamma) \geq S_r$.

If H_r is an exponentially growing stratum, and ρ is a C -quasi-smooth hallway of height r , then we need to develop an understanding of the lengths of components of $\rho_i \cap G_{r-1}$, i.e., we need to study subpaths in G_{r-1} . Intuitively, we will accomplish this by carving out subhallways in G_{r-1} .

Consider a maximal subpath $\gamma \subset G_{r-1}$ of some slice ρ_a , i.e., ρ_a can be expressed as $\alpha\gamma\beta$, and α (resp. β) is either trivial or ends (resp. starts) with a (possibly partial) edge in H_r . We begin the construction of a new hallway ρ' by letting $\rho'_0 = \gamma$.

Now, assume inductively that we have defined the slice ρ'_{i-1} such that ρ'_{i-1} is a maximal subpath of ρ_{a+i-1} in G_{r-1} (we write $\rho_{a+i-1} = \alpha\rho'_{i-1}\beta$), and recall that the slice ρ_{a+i} is obtained by tightening $\mu_{a+i}f(\alpha\rho'_{i-1}\beta)\nu_{a+i}$. We define the notch μ'_i by taking the maximal terminal subpath in G_{r-1} of the path obtained from $\mu_{a+i}f(\alpha)$ by tightening. Similarly, we define the notch ν'_i by tightening the maximal initial subpath in G_{r-1} of the path obtained from $f(\beta)\nu_{a+i}$ by tightening. Observe that tightening $\mu'_i\rho'_{i-1}\nu'_i$ yields a maximal subpath in G_{r-1} of ρ_{a+i} , and that the length of μ'_i and ν'_i is bounded by $C + T_r$. We iterate this

procedure until we reach a point where tightening $\mu'_{i+1}f(\rho'_i)\nu'_{i+1}$ yields a trivial path.

By applying this construction wherever possible, we obtain a fan of $C + T_r$ -quasi-smooth hallways in G_{r-1} . Let \mathcal{M} be the set of maximal elements of this fan. We let \mathcal{M}_1 be the collection of smooth hallways in \mathcal{M} , and we let \mathcal{M}'_2 be the collection of hallways in \mathcal{M} that are not smooth.

Let σ be an element of \mathcal{M}'_2 , and assume that there exists some $0 < i < \mathcal{D}(\sigma)$ such that $\mathcal{L}(\sigma_i) < S_r$. Then we obtain two new hallways σ', σ'' from σ by letting $\sigma'_j = \sigma_j$ for $0 \leq j \leq i$ and $\sigma''_j = \sigma_{i+j}$ for $0 \leq j \leq \mathcal{D}(\sigma) - i$; we may think of this operation as cutting σ along σ_i . We obtain a collection of hallways \mathcal{M}_2 by performing all possible cuts of this kind on all elements of \mathcal{M}'_2 .

If $\sigma \in \mathcal{M}_1 \cup \mathcal{M}_2$, we say that σ *intersects* a slice ρ_i if one of the slices of σ is a subpath of ρ_i . When looking for bounds on the lengths of a slice ρ_i , we need to find bounds on the lengths of slices of hallways σ that intersect ρ_i .

Fix some stratum H_r . We say that the map f satisfies *Condition A_r* if for any $C \geq 0$, there exist computable constants K_r , $K'_r(C)$, and an exponent $d \geq 1$, such that the following two conditions hold:

- If ρ is a smooth hallway in G_r such that the slice ρ_0 starts and ends at fixed vertices, then

$$\mathcal{L}(\rho_i) \leq K_r \mathcal{V}(\rho)$$

for all slices ρ_i .

- If ρ is a C -quasi-smooth hallway in G_r , then

$$\mathcal{L}(\rho_i) \leq K_r \mathcal{V}(\rho) + K'_r(C) \mathcal{D}^d(\rho).$$

If H_r is an exponentially growing stratum, then a hallway of height r is *admissible* if all its slices start and end at fixed vertices or at points in H_r .

Lemma 6.2. *Let H_r be an exponentially growing stratum, and assume that Condition A_{r-1} holds. Then, given some $C \geq 0$, there exist computable constants $C_1, C_2 \geq 1$ with the following property: If ρ is an admissible C -quasi-smooth hallway of height r , then*

$$\mathcal{L}(\rho_i) \leq C_1 \mathcal{V}(\rho) + C_2 \sum_{\substack{\sigma \in \mathcal{M}_2 \\ \sigma \text{ intersects } \rho_i \\ \text{in an } r\text{-significant segment}}} \mathcal{D}^d(\sigma)$$

for every slice ρ_i of ρ .

Proof. Since ρ is admissible, all slices of $\sigma \in \mathcal{M}_1$ start and end at fixed vertices unless σ_0 is contained in a zero stratum, in which case all slices σ_i for $i > 0$ start and end at fixed vertices. Moreover, if σ_0 is contained in a zero stratum, then $\mathcal{L}(\sigma_1) = \mathcal{L}(\sigma_0)$. By Condition A_{r-1} , we have

$$\mathcal{L}(\sigma_i) \leq K_{r-1} \mathcal{V}(\sigma)$$

for all slices σ_i of $\sigma \in \mathcal{M}_1$.

Fix some slice ρ_i of ρ . Using Proposition 4.2 and Condition A_{r-1} , we see that

$$\begin{aligned} \mathcal{L}(\rho_i) &\leq K\mathcal{V}(\rho) + \sum_{\substack{\sigma \in \mathcal{M}_1 \\ \sigma \text{ intersects } \rho_i}} K_{r-1}\mathcal{V}(\sigma) \\ &\quad + \sum_{\substack{\sigma \in \mathcal{M}_2 \\ \sigma \text{ intersects } \rho_i}} (K_{r-1}\mathcal{V}(\sigma) + K'_{r-1}(C + T_r)\mathcal{D}^d(\sigma)). \end{aligned}$$

Consider some $\sigma \in \mathcal{M}_1$ that intersects ρ_i . If the initial slice of σ is not visible in ρ , then, as we noted before, its length is bounded by T_r . Similarly, if the terminal slice of σ is not visible in ρ , then its length is also bounded by T_r . The number of elements of \mathcal{M}_1 that intersect ρ_i is bounded by $K\mathcal{V}(\rho)$. Putting it all together, we conclude that

$$\sum_{\substack{\sigma \in \mathcal{M}_1 \\ \sigma \text{ intersects } \rho_i}} \mathcal{V}(\sigma) \leq (2KT_r + 1)\mathcal{V}(\rho).$$

Similarly, using the fact that elements of \mathcal{M}_2 are $C + T_r$ -quasi-smooth, and that their initial and terminal slices are either visible in ρ or of length less than S_r , we see that

$$\sum_{\substack{\sigma \in \mathcal{M}_2 \\ \sigma \text{ intersects } \rho_i}} \mathcal{V}(\sigma) \leq (2KS_r + 1)\mathcal{V}(\rho) + 2(C + T_r) \sum_{\substack{\sigma \in \mathcal{M}_2 \\ \sigma \text{ intersects } \rho_i}} \mathcal{D}(\sigma).$$

Since ρ_i contains at most $K\mathcal{V}(\rho)$ subpaths in G_{r-1} , the total contribution of subpaths in G_{r-1} that are not r -significant is bounded by $KS_r\mathcal{V}(\rho)$. Letting $C_1 = K + 2K_{r-1}(K(S_r + T_r) + 1) + KS_r$ and $C_2 = K'_{r-1}(C + T_r) + 2(C + T_r)$, we conclude that

$$\mathcal{L}(\rho_i) \leq C_1\mathcal{V}(\rho) + C_2 \sum_{\substack{\sigma \in \mathcal{M}_2 \\ \sigma \text{ intersects } \rho_i \\ \text{in an } r\text{-significant segment}}} \mathcal{D}^d(\sigma).$$

□

Lemma 6.2 shows that from now on, we may focus on the polynomial contribution of nonsmooth hallways in G_{r-1} that intersect a given slice ρ_i in an r -significant subpath. In particular, if the initial slice ρ_0 happens to be an r -legal path, then

$$\mathcal{L}(\rho_i) \leq C_1\mathcal{V}(\rho)$$

for all slices ρ_i since \mathcal{M}_2 is empty in this case.

Lemma 6.3. *Let H_r be an exponentially growing stratum, and assume that Condition A_{r-1} holds. Given some $C > 0$, there exist computable constants C_1, C_2 with the following property: If ρ is an admissible C -quasi-smooth hallway*

of height r , such that for every slice ρ_i except possibly the last one, $f_{\#}(\rho_i)$ does not contain a legal segment of length at least L , then

$$\mathcal{L}(\rho_i) \leq C_1 \mathcal{V}(\rho) + C_2 \mathcal{D}^{d+1}(\rho)$$

for all slices ρ_i .

Proof. By Lemma 6.2, we may restrict our attention to elements of \mathcal{M}_2 that intersect a given slice ρ_i in an r -significant subpath. Let

$$D = \sum_{\substack{\sigma \in \mathcal{M}_2 \\ \sigma \text{ intersects } \rho_i \\ \text{in an } r\text{-significant segment}}} \mathcal{D}^d(\sigma).$$

We first claim that the number of r -significant subpaths in G_{r-1} in a slice ρ_i is bounded by $N(\rho_i)$. By choice of S_r , an r -significant subpath in G_{r-1} will not cancel completely when $f(\rho_i)$ is tightened to $f_{\#}(\rho_i)$.

If there were two such subpaths in one legal segment of ρ_i , then there would be a legal segment in H_r in between. Since we assumed that $\mathcal{L}(f(E) \cap H_r) \geq L$ for each edge in H_r , the r -length of the image of this legal segment is at least L , which means that the slice ρ_{i+1} contains a legal segment of length at least L , contradicting our assumption. This proves the claim if H_r does not support a closed Nielsen path, as in this case, the number of legal segments in ρ_i equals $N(\rho_i)$.

If H_r supports a closed Nielsen path, then a legal segment of ρ_i that is adjacent to an illegal turn contained in a Nielsen subpath of ρ_i cannot contain an r -significant subpath in G_{r-1} . If such a segment contained an r -significant subpath in G_{r-1} , then $f_{\#}(\rho_i)$ would contain a legal segment of r -length L because both the initial and terminal partial edge of the Nielsen path of H_r map to legal segments of r -length at least L . This implies that the number of r -significant subpaths in G_{r-1} is bounded by $N(\rho_i)$.

Now, fix some slice ρ_i . We make the worst-case assumption that every legal segment of ρ that is not adjacent to an illegal turn contained in a Nielsen subpath contains an r -significant subpath in G_{r-1} that is a slice of a hallway $\sigma \in \mathcal{M}_2$ of duration $j \geq i$. The number of such hallways whose duration is a given number $j \geq i$ is bounded by $N(\rho_j) + 1$. We conclude that

$$D \leq \sum_{j=i}^{\mathcal{D}(\rho)} N(\rho_j) j^d.$$

Choosing λ according to Lemma 2.5, we conclude that $N(\rho_{i+1}) \leq \lambda^{-1} N(\rho_i) + 1 + 2C$, as ρ is C -quasi-smooth. This implies, inductively, that

$$N(\rho_i) \leq \lambda^{-i} N(\rho_0) + 2(1 + C) \sum_{j=0}^{i-1} \lambda^{-j} \leq \lambda^{-i} N(\rho_0) + \frac{\lambda}{\lambda - 1} (1 + 2C).$$

We choose some $B \geq \sum_{j=0}^{\infty} \lambda^{-j} j^d$, and we conclude that

$$D \leq \sum_{j=0}^{\mathcal{D}(\rho)} N(\rho_j) j^d \leq B\mathcal{V}(\rho) + \frac{\lambda}{\lambda-1} (1+2C)\mathcal{D}^{d+1}(\rho),$$

since $N(\rho_0) \leq \mathcal{V}(\rho)$.

If C'_1, C'_2 are the constants from Lemma 6.2, then the lemma holds with $C_1 = C'_1 + C'_2 B$ and $C_2 = \frac{\lambda}{\lambda-1} (1+2C)C'_2$. \square

Let H_r be an exponentially growing stratum, and let N_0 be the constant from Lemma 2.5. We say that an admissible smooth hallway ρ of height r has *Property B* if for all slices ρ_i , ρ_i contains no long r -legal segment, or $N(\rho_{i-1}) < N_0$.

Lemma 6.4. *Let H_r be an exponentially growing stratum, and assume that Condition A_{r-1} holds. Let N_0 be the constant from Lemma 2.5. There exist computable constants C_1, C_2 with the following property: If ρ is an admissible smooth hallway of height r that satisfies Property B, then*

$$\mathcal{L}(\rho_i) \leq C_1 \mathcal{V}(\rho) + C_2 \mathcal{D}^{d+1}(\rho)$$

for all slices ρ_i .

Proof. If no slice of ρ contains a long legal segment, then the claim follows from Lemma 6.3. Otherwise, let i_0 be the smallest index for which ρ_{i_0} contains a long legal segment. By choice of i_0 , ρ_{i_0-1} does not contain a long legal segment, and by hypothesis, we have $N(\rho_{i_0-1}) < N_0$. If $i < i_0$, then, choosing D as in the proof of Lemma 6.3, we conclude that

$$\begin{aligned} D &\leq \left(\sum_{j=0}^{i_0-1} N(\rho_j) j^d \right) + N_0 \mathcal{D}^d(\rho) \\ &\leq B\mathcal{V}(\rho) + N_0 \mathcal{D}^{d+1}(\rho), \end{aligned}$$

so that the lemma holds for all ρ_i with $i < i_0$.

For $i \geq i_0$, ρ_i splits as a concatenation of long r -legal paths and subpaths that contain illegal turns and no long legal subpaths. Each slice may, conceivably, contain slices of $N(\rho_{i_0-1}) < N_0$ hallways of duration $\mathcal{D}(\rho)$. The polynomial contribution of these hallways is bounded by $N_0 \mathcal{D}^d(\rho)$.

In addition, the number of short legal segments around illegal turns is at most $2N_0$. Each of them contains not more than one r -significant subpath in G_{r-1} , belonging to a hallway of duration at most $\mathcal{D}(\rho) - i_0$. The polynomial contribution of these paths is bounded by $2N_0 (\mathcal{D}(\rho) - i_0)^d$.

Now, since ρ_{i_0} contains a long legal segment, the length of $\rho_{\mathcal{D}(\rho)} = f_{\#}^{\mathcal{D}(\rho)-i_0}(\rho_{i_0})$ is at least $\lambda_r^{\mathcal{D}(\rho)-i_0}$. We can easily find some $B' > 0$ such that $B' \lambda_r^k \geq 2N_0 k^d$

for all $k \geq 0$. We conclude that for the sum of all polynomial contributions in ρ_i , we have

$$2N_0(\mathcal{D}(\rho) - i_0)^d + N_0\mathcal{D}^d(\rho) \leq B'\mathcal{V}(\rho) + N_0\mathcal{D}^d(\rho),$$

which completes the proof of the lemma. \square

The remaining two lemmas deal with arbitrary smooth hallways of height r as well as quasi-smooth hallways by essentially decomposing them into pieces of the kind that we analyzed in the previous lemmas.

Lemma 6.5. *Let H_r be an exponentially growing stratum, and assume that Condition A_{r-1} holds. Then there exist computable constants C_1, C_2 with the following property: If ρ is a smooth admissible hallway of height r , then*

$$\mathcal{L}(\rho_i) \leq C_1\mathcal{V}(\rho) + C_2\mathcal{D}^{d+1}(\rho)$$

for all slices ρ_i .

Proof. Let λ_-, N_0 be the constants from Lemma 2.5. As in the proof of Lemma 6.1, we let $\lambda = \sqrt{\min\{\lambda_-, \lambda_r\}}$, and we remark that for $0 \leq j \leq k$, we have $\lambda_-^j + \lambda_r^{k-j} \geq \lambda^k$. This basic estimate will be crucial in the proof of this lemma. We choose some $B > 0$ such that $B\lambda^k > k^{d+1}$ for all $k \geq 0$.

Let C'_1, C'_2 be the maximum of the corresponding constants from the previous lemmas. We will see that the lemma holds with $C_1 = C'_1 + 3BC'_2$ and $C_2 = C'_2$.

We first observe that if ρ satisfies Property B , then the lemma follows from Lemma 6.4. If ρ_0 contains long legal segments, we can split ρ_0 into long r -legal subpaths and neighborhoods of illegal turns (i.e., illegal turns surrounded by legal paths whose length is at most $\frac{C_r}{2}$). Split ρ_0 as $\rho_0 = \alpha_{0;1}\beta_{0;1}\alpha_{0;2} \cdots \alpha_{0;m}\beta_{0;m}$, where all subpaths $\alpha_{0;i}$ are long legal segments, and all subpaths $\beta_{0;i}$ are neighborhoods of illegal turns. Such a decomposition of ρ_0 induces a decomposition of ρ into hallways, and we can choose the decomposition such that all resulting pieces are admissible, and that the legal segments are as long as possible, subject to admissibility. We write $\alpha_{j;i} = f_{\#}^j(\alpha_{0;i})$ and $\beta_{j;i} = f_{\#}^j(\beta_{0;i})$.

Let $k = \mathcal{D}(\rho)$. For each long legal subpath $\alpha_{0;i}$, Lemma 6.2 yields that $\mathcal{L}(\alpha_{j;i}) \leq C_1(\mathcal{L}(\alpha_{0;i}) + \mathcal{L}(\alpha_{k;i}))$, for all $0 \leq j \leq k$. Since $\alpha_{0;i}$ is a long legal segment, we have $\mathcal{L}(\alpha_{k;i}) \geq \lambda_r^k \geq \lambda^k$.

If the hallway defined by $\beta_{0;i}$ satisfies Property B , then we have $\mathcal{L}(\beta_{j;i}) \leq C'_1(\mathcal{L}(\beta_{0;i}) + \mathcal{L}(\beta_{k;i})) + C'_2k^{d+1}$, and we have $k^{d+1} \leq B\mathcal{L}(\alpha_{k;i})$, hence

$$\mathcal{L}(\beta_{j;i}) \leq C'_1(\mathcal{L}(\beta_{0;i}) + \mathcal{L}(\beta_{k;i})) + BC'_2\mathcal{L}(\alpha_{k;i}),$$

i.e., we can find a legal segment adjacent to $\beta_{0;i}$ whose contribution to the visible edges of ρ dominates the possible polynomial contribution of $\beta_{0;i}$. This takes care of the long legal segments in ρ_0 as well as the subpaths that satisfy Property B . Hence, we only need to deal with those paths that do not satisfy Property B . Assume that for some $0 \leq i \leq m$, $\beta_{0;i}$ is one of them.

Then there exists some j_0 such that $\beta_{j_0;i}$ contains a long legal segment, but $\beta_{j_0-1;i}$ does not, and $N(\beta_{j_0-1;i}) \geq N_0$.

As before, we split $\beta_{j_0;i}$ into long legal segments and neighborhoods of illegal turns, obtaining a decomposition $\beta_{j_0,i} = \alpha_{j_0;i,0}\beta_{j_0;i,0} \cdots \alpha_{j_0;i,m}\beta_{j_0;i,m}$, where $\alpha_{j_0;i,k}$ are r -legal subpaths, and $\beta_{j_0;i,k}$ are neighborhoods of illegal turns. We can find splittings $\beta_{j;i} = \alpha_{j;i,0}\beta_{j;i,0} \cdots \alpha_{j;i,m}\beta_{j;i,m}$ for all $0 \leq j \leq k$, such that $f_{\#}(\alpha_{j;i,k}) = \alpha_{j+1;i,k}$ and $f_{\#}(\beta_{j;i,k}) = \beta_{j+1;i,k}$. We may choose those splitting such that the resulting pieces are admissible, and such that the legal segments $\alpha_{j_0;i,k}$ are as long as possible, subject to admissibility.

Now, fix on one subpath $\alpha_{j_0;i,k}$. If N is the number of r -significant subpaths in G_{r-1} in $\alpha_{j_0;i,k}$, then $\alpha_{j_0-1;i,k}$ contains at least N legal segments containing r -significant subpaths in G_{r-1} . By Lemma 2.5, we have $\mathcal{L}(\beta_{0;i}) \geq N(\beta_{j_0;i}) \geq \lambda_-^{j_0-1} N(\beta_{j_0-1;i})$, so that we can find $\lambda_-^{j_0-1} N$ illegal turns in $\beta_{0;i}$, and we can find $\lambda_r^{k-j_0}$ edges in $\beta_{k,i}$. Using our earlier estimate, we see that $(\lambda_-^{j_0-1} + \lambda_r^{k-j_0})N \geq \lambda_-^{-1} \lambda^k N$.

The polynomial contribution of the r -significant subpaths in G_{r-1} of $\alpha_{j_0;i,k}$ is bounded by $K'_{r-1}(T_r)Nk^{d+1} \leq BK'_{r-1}(T_r)N\lambda^k$, i.e., it is dominated by corresponding visible edges.

This leaves us to deal with the adjacent subpaths $\beta_{j_0;i,k}$ and $\beta_{j_0;i,k-1}$. If $\beta_{0;i,k}$ satisfies Property B , then its polynomial contribution is bounded by $C'_2 k^{d+1}$, which in turn is bounded by $BC'_2 \lambda^k$.

This takes care of the legal segments $\alpha_{j_0;i,k}$ as well as those neighborhoods of illegal turns that satisfy Property B . We apply the previous reasoning to the remaining paths $\beta_{j_0;i,k}$, completing the proof of the lemma. \square

Lemma 6.6. *Let H_r be an exponentially growing stratum, and assume that Condition A_{r-1} holds. Given some $C > 0$, there exist computable constants C_1, C_2 with the following property: If ρ is an admissible C -quasi-smooth hallway of height r , then*

$$\mathcal{L}(\rho_i) \leq C_1 \mathcal{V}(\rho) + C_2 \mathcal{D}^{d+3}(\rho)$$

for all slices ρ_i .

Proof. The idea of this proof is to decompose the hallway ρ into pieces that are either smooth or C -quasi-smooth satisfying the hypothesis of Lemma 6.3.

In order to find this decomposition, we introduce *trajectories* of points in H_r . This definition may be affected by the choices made when tightening (Remark 1.1). In order to avoid ambiguities, for each index $1 \leq i < D(\rho)$, we fix a sequence of elementary cancellations that turn $\mu_i \rho_{i-1} \nu_i$ into ρ_i .

If p is a point in $\rho_i \cap H_r$, we consider its image $f(p)$ in $f(\rho_i)$. We say that p *survives* if $f(p)$ is contained in H_r and if $f(p)$ is not contained in an edge that cancels when $f(\rho_i)$ is tightened to $f_{\#}(\rho_i)$. If p survives, then $f(p)$ is contained in ρ_{i+1} , or it is contained in the parts of $f_{\#}(\rho_i)$ that cancel when $\mu_{i+1} f_{\#}(\rho_i) \nu_{i+1}$ is tightened to ρ_{i+1} .

Thinking of the hallway ρ as spanning a (possibly singular) disk, we draw a line segment (in this disk) from the surviving points in each slice to their images. If p is a point in a visible edge such that p and all its images survive, then p defines a line starting and ending in visible edges, called the *trajectory* of

p . The trajectories of two points need not be disjoint, but that does not concern us here.

We say that two trajectories are *parallel* if their initial points are both contained in ρ_0 or both contained in the same notch, and if their terminal points are both contained in $\rho_{\mathcal{D}(\rho)}$ or both contained in the same notch. The crucial observation is that equivalence classes of parallel trajectories are closed subsets of the disk spanned by ρ , so that in every equivalence class, we can find trajectories of two points p_1, p_2 that are extremal in the following sense: If p is a point whose trajectory is parallel to those of p_1 and p_2 , then p is located between p_1 and p_2 .

We now cut ρ along the extremal trajectories of all equivalence classes of parallel trajectories, obtaining pieces that are either smooth or C -quasi-smooth. Moreover, all the resulting pieces are admissible. Let \mathcal{M}_1 be the collection of smooth pieces and \mathcal{M}_2 the collection of pieces that are not smooth. Note that $\mathcal{V}(\mathcal{M}_1) + \mathcal{V}(\mathcal{M}_2) = \mathcal{V}(\rho)$.

We now claim that all elements of \mathcal{M}_2 satisfy the hypothesis of Lemma 6.6. Suppose otherwise, i.e., there exists some $\sigma \in \mathcal{M}_2$ such that for some slice σ_i , $f_{\#}(\sigma_i)$ contains a legal segment of length at least L . Within the interior of this legal segment, we can find some point p such that all images of p survive in subsequent slices. Since p is the image of surviving points, we obtain a trajectory along which we can cut σ , contradicting the fact that we obtained σ by cutting ρ along extremal trajectories.

By Lemma 6.5, there are constants C'_1, C'_2 such that for every $\sigma \in \mathcal{M}_1$ and every slice σ_i of σ , we have

$$\mathcal{L}(\sigma_i) \leq C'_1 \mathcal{V}(\sigma) + C'_2 \mathcal{D}^{d+1}(\sigma),$$

and by Lemma 6.6, there are constants C''_1, C''_2 such that

$$\mathcal{L}(\sigma_i) \leq C''_1 \mathcal{V}(\sigma) + C''_2 \mathcal{D}^{d+1}(\sigma)$$

for every slice σ_i of every $\sigma \in \mathcal{M}_2$.

There are at most $2(\mathcal{D}(\rho) - 1)$ notches, so that the number of equivalence classes of parallel trajectories is bounded by $(2(\mathcal{D}(\rho) - 1) + 1)^2$ (another extremely crude estimate, but it'll do). Since we cut along no more than two trajectories per equivalence class, we obtain no more than

$$2(2(\mathcal{D}(\rho) - 1) + 1)^2 + 1 \leq 8\mathcal{D}^2(\rho)$$

pieces. Letting $C_1 = \max\{C'_1, C''_1\}$ and $C_2 = 8 \max\{C'_2, C''_2\}$, we conclude that

$$\mathcal{L}(\rho_i) \leq C_1 \mathcal{V}(\rho) + C_2 \mathcal{D}^{d+3}(\rho)$$

for all slices of ρ . □

We now have all the ingredients that we need to prove Theorem 1.9.

Proof of Theorem 1.9. We first show that Condition A_r holds for all strata H_r . This implies, in particular, that the second statement of Theorem 1.9 holds for paths starting and ending at fixed vertices. If ρ is a path starting and ending at arbitrary vertices, then Theorem 1.4, Part 2 yields that $f_{\#}(\rho)$ starts and ends at *fixed* vertices, so that, in fact, the second statement of Theorem 1.9 follows from Condition A_r in this case as well.

We note that Condition A_0 holds trivially, and we assume inductively that Condition A_{r-1} holds for some r . We want to prove Condition A_r .

Assume that H_r is an exponentially growing stratum, and let ρ be a smooth hallway of height r such that ρ_0 starts and ends at fixed vertices. If ρ_0 is a concatenation of Nielsen paths of height r and paths in G_{r-1} , then we can split ρ_0 at the endpoints of its subpaths in G_{r-1} , and Condition A_{r-1} completes the proof. We now assume that ρ_0 is not a concatenation of Nielsen paths and paths in G_{r-1} .

By Lemma 6.5, we have constants C_1, C_2 such that

$$\mathcal{L}(\rho_i) \leq C_1 \mathcal{V}(\rho) + C_2 \mathcal{D}^{d+1}(\rho)$$

for all slices ρ_i . Moreover, by Lemma 6.1, there exists some $C > 0$ and $\lambda > 1$, independently of ρ , such that

$$\mathcal{V}(\rho) \geq C \lambda^{\mathcal{D}(\rho)}.$$

We can easily find some constant B such that $BC\lambda^k \geq C_2 k^{d+1}$ for all $k \geq 0$. Now the first part of Condition A_r follows, with $K_r = C_1 + B$. Lemma 6.6 yields the second part of Condition A_r , so that Condition A_r holds.

We now assume that H_r is a polynomially growing stratum. Because of Proposition 5.5, we only need to consider the following situation: Either H_r is fast, or H_r is truly polynomial, but ρ contains fast polynomial edges or non-Nielsen subpaths in exponentially growing strata.

In order to see that the second part of Condition A_r holds for a C -quasismooth hallway ρ of height r , we apply cutting and sawtooth constructions to ρ , obtaining a collection of $(C + |u_r|)$ -quasismooth hallways of height $r - 1$ or less, so that the second part of Condition A_r immediately follows from the second part of Condition A_{r-1} .

Now, given a smooth hallway ρ of height r , we apply cutting and sawtooth constructions again, obtaining a collection of $2|u_r|$ -quasismooth hallways. For each slice ρ_i , the second part of Condition A_{r-1} yields a polynomial bound on the number of edges marked by linear strata (Definition 4.1). Now, since either H_r is fast or ρ contains fast polynomial edges or non-Nielsen subpaths in exponentially growing strata, Lemma 6.1 provides an exponential lower bound for the number of visible edges. As before, the exponential lower bound for visible edges easily dominates the polynomial lower bound for edges marked by linear strata, which completes the proof of Condition A_r .

Finally, in order to prove the first part of Theorem 1.9, we need to understand the dynamics of circuits. Let σ be a circuit of height r . If H_r is a polynomially growing stratum, then Remark 1.6 yields that σ splits, at fixed vertices, into

basic paths of height r and paths in G_{r-1} , so that Condition A_r proves the claim.

Assume that H_r is an exponentially growing stratum. If σ is a concatenation of Nielsen paths of height r and paths in G_{r-1} , then we can split σ at the endpoints of its subpaths in G_{r-1} , so that Condition A_{r-1} completes the proof in this case. We now assume that σ is not a concatenation of Nielsen paths and subpaths in G_{r-1} . Then σ splits at a point p in H_r , so that we may interpret σ as a path starting and ending at v . Let ρ be a smooth hallway with $\rho_0 = \sigma$. Then, by Lemma 6.5, we can find constants C_1, C_2 such that

$$\mathcal{L}(\rho_i) \leq C_1 \mathcal{V}(\rho) + C_2 \mathcal{D}^d(\rho)$$

for all slices ρ_i . Moreover, by Lemma 6.1, we can find constants C, λ such that

$$\mathcal{V}(\rho) \geq C \lambda^{\mathcal{D}(\rho)}.$$

As before, we find some constant B such that $BC\lambda^k \geq C_2 k^d$ for all $k \geq 0$, so that the first statement of Theorem 1.9 holds with $K_r = C_1 + B$.

Finally, if ρ_0 is a Nielsen path of height r , then there is nothing to show. This completes the proof. \square

References

- [BF92] M. Bestvina and M. Feighn. A combination theorem for negatively curved groups. *J. Differential Geom.*, 35(1):85–101, 1992.
- [BFH97] M. Bestvina, M. Feighn, and M. Handel. Laminations, trees, and irreducible automorphisms of free groups. *Geom. Funct. Anal.*, 7(2):215–244, 1997.
- [BFH00] Mladen Bestvina, Mark Feighn, and Michael Handel. The Tits alternative for $\text{Out}(F_n)$. I. Dynamics of exponentially-growing automorphisms. *Ann. of Math. (2)*, 151(2):517–623, 2000.
- [BG] Martin R. Bridson and Daniel P. Groves. The quadratic isoperimetric inequality for mapping tori of free group automorphisms II: The general case. arXiv:math.GR/0610332.
- [BH92] Mladen Bestvina and Michael Handel. Train tracks and automorphisms of free groups. *Ann. of Math. (2)*, 135(1):1–51, 1992.
- [Bri00] Peter Brinkmann. Hyperbolic automorphisms of free groups. *Geom. Funct. Anal.*, 10(5):1071–1089, 2000. arXiv:math.GR/9906008.
- [Bri07] Peter Brinkmann. Detecting orbits of free group automorphisms, 2007. preprint.
- [Coo87] Daryl Cooper. Automorphisms of free groups have finitely generated fixed point sets. *J. Algebra*, 111(2):453–456, 1987.

- [DV96] Warren Dicks and Enric Ventura. *The group fixed by a family of injective endomorphisms of a free group*. American Mathematical Society, Providence, RI, 1996.
- [Gan59] F. R. Gantmacher. *The theory of matrices. Vols. 1, 2*. Chelsea Publishing Co., New York, 1959. Translated by K. A. Hirsch.
- [Ger94] S. M. Gersten. The automorphism group of a free group is not a $\text{CAT}(0)$ group. *Proc. Amer. Math. Soc.*, 121(4):999–1002, 1994.
- [Mac00] N. Macura. Quadratic isoperimetric inequality for mapping tori of polynomially growing automorphisms of free groups. *Geom. Funct. Anal.*, 10(4):874–901, 2000.

DEPARTMENT OF MATHEMATICS
 THE CITY COLLEGE OF CUNY
 NEW YORK, NY 10031
 U.S.A.

E-mail: brinkman@sci.ccny.cuny.edu